EXHIBIT 15

UWEE Faculty

Jeff Bilmes
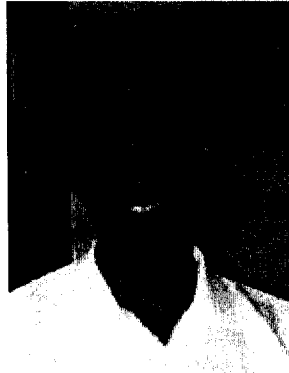
Jeff Bilmes
Assistant Professor
Signal and Image Processing
418 EE/CSE
Box 352500
University of Washington
Seattle, WA 98195

Personal homepage

Phone: (206) 543-2150
E-mail: bilmes@ee.washington.edu

University of California Berkeley, 1999 Ph.D.
Massachusetts Institute of Technology, 1993 M.S.
University of California Berkeley, 1989 B.S.

[Biosketch] [Honors] [Research Interests] [More Information]

Biosketch

Jeff A. Bilmes joined the University of Washington Department of
Electrical Engineering faculty in the fall of 1999. He received a B.S.
degree in Electrical Engineering and Computer Science from U.C.
Berkeley in 1989, a S.M. degree from MIT in 1993, and a Ph.D. in
Computer Science from U.C. Berkeley in 1999. For many years, he als
worked and performed research at the International Computer Science
Institute in Berkeley, CA. He is the author of about 45 journal and
conference papers on topics ranging from speech and pattern
recognition to parallel programming language design to high-
performance software coding techniques.

He is an associate editor of the IEEE Transactions on Multimedia, and
also an adjunct professor in the University of Washington Department
Linguistics

and in Computer Science and Engineering

In winter 2003, he co-leading IEEE ASRU2003, the IEEE Workshop on Automatic Speech Recognition and Understanding.

In summer 2001, he lead the JHU CSLP workshop on discriminatively structured graphical models for speech recognition.

Research Interests

His primary interests lie in statistical modeling (particularly graphical approaches) and signal processing for pattern classification, speech recognition, language processing, and audio processing. He is particularly interested in discovering ways to discriminatively structure such models for classification problems. He also has strong interests in speech-based human-computer interfaces, general machine learning, the statistical properties of natural objects and natural scenes, information theory and its relation to natural computation by humans and pattern recognition by machines, and computational music processing (such as human timing subtleties). He is also quite interested in high performance computing systems, computer architecture, and software techniques to reduce power consumption.

Honors

2001 NSF CAREER Award recipient
2001 CRA Digital Government Fellow
1998 U.C. Berkeley EECS Dept. Samuel Silver Scholarship Award

Jeff is a member of the IEEE Computer Society, Signal Processing Society, Information Theory Society, Systems, Man, and Cybernetics Society, and also the ACM Computer Society.

For More Information

Much more information is available on my web page at the deparment Signal, Speech, and Language Interpretation laboratory.

Computer Speech and Langauge : Special Issue

I am guest editing, along with Martin Russel, a speecial issue of Computer Speech and Language on new computational paradigms for acoustic modeling in speech recognition. The deadline for submission i October 31st, and the PDF call for papers is here.

College of Engineering | University of Washington | Site Search | Questions or Problems? | Giving

Jeff A. Bilmes

Assistant Professor (Associate as of Oct.'05) , Electrical Engineering
Adjunct Professor in Computer Science and Engineering, and in Linguistics.
University of Washington, Seattle
Signal, Speech, and Language Interpretation Laboratory
418 EE/CS Bldg, Box 352500
Seattle, Washington 98195-2500
(206) 221-5236 fax: (206)543-3842
bilmes AT ee DOT washington DOT edu
Departmental Web Page: http://www.ee.washington.edu/people/faculty/bilmes

Download vcard

If you need to contact me, you may also leave a message with my administrative assistant Kathy Burch
either by email at kathy@ssli.ee.washington.edu or by telephone at 206-616-1763

General Information

Since August 1999, I have been an assistant professor at the Department of Electrical Engineering at the University of Washington, Seattle Washington. I am also an adjunct professor in computer science and the department of Linguistics. I am also a co-founder of the Signal, Speech, and Language Interpretation Laboratory here in the department. I received my Ph.D. from the Computer Science Division of the department of Electrical Engineering and Computer Science, University of California in Berkeley. I was also a researcher at the International Computer Science Instiute, and a member of the Realization group there.

My primary interests lie in statistical modeling (particularly graphical model approaches) and signal processing for pattern classification, speech recognition, language processing, data fusion and data fission, and audio processing. I am particularly interested in discriminatively structured graphical models for classification. I also have strong interests in speech-based human-computer interfaces, general machine learning, the statistical properties of natural objects and natural scenes, information theory and its relation to natural computation by humans and pattern recognition by machines, and computational music processing (such as human timing subtleties). I am also quite interested in high performance computing systems, computer architecture, and software techniques to reduce power consumption.

In the summer 2001, I was a leader at the JHU CSLP workshop on discriminatively structured graphical models for speech recognition based on the ideas presented in my 2000 UAI paper

Brief Curriculum Vitae

1989	University of California Berkeley
	Department of Electrical Engineering and Computer Science
	Computer Science Division B.S. in Electrical Engineering and Computer Science
1989-1991	International Computer Science Institute
	Berkeley, CA Research Programmer
1991-1993	Massachusetts Institute of Technology
	Master of Science
1993-1999	University of California Berkeley
	Department of Electrical Engineering and Computer Science
	Computer Science Division Ph.D. in Computer Science, May 1999.
1999-	University of Washington Seattle
	Department of Electrical Engineering
	Assistant Professor

Publications

- Refereed Conference Proceedings
- Journal Publications and Book Chapters
- Tutorials
- Technical Reports
- Technical Documentation

Refereed Conference Proceedings

- Yi Li, Linda Shapiro, and Jeff Bilmes **A Generative/Discriminative Learning Algorithm for Object Recognition** (to appear) *10th IEEE Conference on Computer Vision (ICCV 2005)* , Beijing, China 2005

- Franz Pernkopf and Jeff Bilmes **Discriminative versus Generative Parameter and Structure Learning of Bayesian Network Classifiers** (gzipped ps or pdf) (to appear) *International Conference on Machine Learning* , Bonn, Germany, 2005

- Mukund Narasimhan and Jeff Bilmes **A Supermodular-Submodular Procedure with Applications to Discriminative Structure Learning.** (gzipped ps or pdf) (to appear) *21st Conference on Uncertainty in Artificial Intelligence (UAI05)* , Edinburgh, Scotland, 2005

- Xiao Li, Jeff Bilmes, and Jon Malkin **Maximum Margin Learning and Adaptation of MLP Classifers** (gzipped ps or pdf) (to appear) *9th European Conference on Speech Communication and Technology (Eurospeech'05)* , Lisbon, Portugal, 2005

- Chris Bartels, Kevin Duh, Jeff Bilmes, Katrin Kirchhoff, and Simon King **Genetic Triangulation of Graphical Models for Speech and Language Processing** (gzipped ps or pdf) (to appear) *9th European Conference on Speech Communication and Technology (Eurospeech'05)* , Lisbon, Portugal, 2005

- Amar Subramanya, Jeff Bilmes, and Chia-Ping Chen **Focused Word Segmentation for ASR** (gzipped ps or pdf) (to appear) *9th European Conference on Speech Communication and Technology (Eurospeech'05)* , Lisbon, Portugal, 2005

- Simon King, Chris Bartels, and Jeff Bilmes **SVitchboard 1: Small Vocabulary Tasks from Switchboard 1** (gzipped ps or pdf) (to appear) *9th European Conference on Speech Communication and Technology (Eurospeech'05)* , Lisbon, Portugal, 2005

- Karim Filali and Jeff Bilmes **A Dynamic Bayesian Framework to Model Context and Memory in Edit Distance Learning: An Application to Pronunciation Classification** (gzipped ps or pdf) (to appear) *Proceedings of the Association for Computational Linguistics (ACL)*, University of Michigan, Ann Arbor, 2005

- Gang Ji and Jeff Bilmes **Dialog Act Tagging using Graphical Models** (gzipped ps or pdf) *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Philadelphia, PA, March 2005

- Xin Lei, Gang Ji, Tim Ng, Jeff Bilmes, and Mari Ostendorf **DBN-Based Multi-stream Models for Mandarin Toneme Recognition** (gzipped ps or pdf) *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Philadelphia, PA, March 2005

- Jon Malkin, Xiao Li, and Jeff Bilmes **A Graphical Model for Formant Tracking** (gzipped ps or pdf) *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Philadelphia, PA, March 2005

- Chia-Ping Chen, Jeff Bilmes, Dan Ellis **Speech Feature Smoothing for Robust ASR**

(gzipped ps or pdf) *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* Philadelphia, PA, March 2005

- Mukund Narasimhan, Jeff Bilmes **Optimal Sub-graphical Models** (gzipped ps or pdf) *Neural Information Processing Systems (NIPS), 2004* , Vancover, Canada, Dec 2004
- Xiao Li, Jonathan Malkin, and Jeff A. Bilmes **A Graphical Model Approach to Pitch Tracking** (gzipped ps or pdf) *International Conference on Spoken Language Processing (ICSLP) 2004* , October 2004.
- Mukund Narasimhan and Jeff A. Bilmes **PAC-learning bounded tree-width Graphical Models** (gzipped ps or pdf) *20th Conference on Uncertainty in Artificial Intelligence (UAI)* , July 2004.
- Yi Li, Jeff A. Bilmes, and Linda G. Shapiro **Object Class Recognition using Images of Abstract Regions** (gzipped ps or pdf) *International Conference on Pattern Recognition,* August 2004. Cambridge, UK
- Gang Ji and Jeff Bilmes. **Multi-Speaker Language Modeling** (gzipped ps or pdf) *Human Language Technology Conference,* May 2004. Boston, MA
- John N. Gowdy, Amarnag Subramanya, Chris Bartels, and Jeff Bilmes. **DBN-Based Multi-Stream Models for Audio-Visual Speech Recognition (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* May 2004. Montreal, Canada
- Xiao Li, Jonathan Malkin, and Jeff Bilmes. **Codebook Design for ASR Systems using Custom Arithmetic units (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* May 2004. Montreal, Canada
- Jonathan Malkin, Xiao Li, and Jeff Bilmes. **Custom Arithmetic for High-Speed, Low-Resource ASR Systems (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* May 2004. Montreal, Canada
- Gang Ji and Jeff Bilmes **Necessary Intransitive Likelihood-Ratio Classifiers** (gzipped ps or pdf) *Neural Information Processing Systems (NIPS), 2003* , 2003. Vancover, Canada, Dec 2003
- Xiao Li and Jeff Bilmes **Feature Pruning in Likelihood Evaluation of HMM-based Speech Recognition** (gzipped ps or pdf) *IEEE Automatic Speech Recognition and Understanding (ASRU), 2003 (to appear)* , 2003. St. Thomas, U.S. Virgin Islands, Dec 2003
- Jeff Bilmes and Chris Bartels **On Triangulating Dynamic Graphical Models** (gzipped ps or pdf) *19th Conference on Uncertainty in Artificial Intelligence (UAI)* , 2003. Acapulco, Mexico.
- Karen Livescu, James Glass, and Jeff Bilmes **Hidden Feature Models for Speech Recognition Using Dynamic Bayesian Networks (gzipped ps or pdf)** *8th European Conference on Speech Communication and Technology (Eurospeech)* , 2003. Geneva, Switzerland.
- Jeff Bilmes and Katrin Kirchhoff **Factored Language Models and Generalized Parallel Backoff (gzipped ps or pdf)** *Human Language Technology Conference* , 2003. Edmonton, Canada.
- Katrin Kirchhoff, Jeff Bilmes, Sourin Das, Nicolae Duta, Melissa Egan, Gang Ji, Feng He, John Henderson, Daben Liu, Mohamed Noamany, Pat Schone, Richard Schwartz, Dimitra Vergyri **Novel Approaches to Arabic Speech Recognition: Report from the 2002 Johns-Hopkins Summer Workshop (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* April 2003. Hong Kong, China

- Yimin Zhang, Qian Diao, Shan Huang, Wei Hu, Chris Bartels, and Jeff Bilmes. **DBN Based Multi-Stream Models for Speech (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* April 2003. Hong Kong, China

- Ozgur Cetin, Harriet Nock, Katrin Kirchhoff, Jeff Bilmes, and Mari Ostendorf **The 2001 GMTK-Based SPINE ASR System (gzipped ps or pdf)** *International Conference on Spoken Language Processing (ICSLP) 2002* Denver, Colorado

- Karim Filali, Xiao Li, and Jeff Bilmes **Data-Driven Vector Clustering for Low-Memory Footprint ASR (gzipped ps or pdf)** *International Conference on Spoken Language Processing (ICSLP) 2002* Denver, Colorado

- Chia-Ping Chen, Karim Filali, and Jeff Bilmes **Frontend Post-Processing and Backend Model Enhancement on the Aurora 2.0/3.0 Databases (gzipped ps or pdf)** *International Conference on Spoken Language Processing (ICSLP) 2002* Denver, Colorado (also see the MVA Web site)

- Chia-Ping Chen, Jeff Bilmes, and Katrin Kirchhoff **Low-Resource Noise-Robust Feature Post-Processing on Aurora 2.0 (gzipped ps or pdf)** *International Conference on Spoken Language Processing (ICSLP) 2002* Denver, Colorado (also see the MVA Web site)

- J. Bilmes and G. Zweig **The Graphical Models Toolkit: An Open Source Software System for Speech and Time-Series Processing (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* June 2002. Orlando Florida

- G. Zweig, J. Bilmes, T. Richardson, K. Filali, K. Livescu, P. Xu, K. Jackson, Y. Brandman, E. Sandness, E. Holtz, J. Torres, B. Byrne **Structurally Discriminative Graphical Models for Automatic Speech Recognition: Results from the 2001 Johns Hopkins Summer Workshop (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* June 2002. Orlando Florida

- K. Kirchhoff, S. Parandekar, and J. Bilmes **Mixed-Memory Markov Models for Automatic Language Identification (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* June 2002. Orlando Florida

- I. Bulyko, M. Ostendorf, and J. Bilmes **Robust Splicing Costs and Efficient Search with BMM Models for Concatenative Speech Synthesis (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing,* June 2002.

- J. Bilmes, G. Ji, and M. Meila **Intransitive Likelihood-Ratio Classifiers (gzipped ps or pdf)** *NIPS'2001* Vancouver Canada, Dec 2001

- R. Vuduc, J. Demmel, and J. Bilmes **On Statistical Models in Automatic Tuning (gzipped ps or pdf)** *The International Conference on Computational Science: Architecture-Specific Automatic Performance Tuning* San Francisco, May 2001

- R. Vuduc, J. Bilmes, and J. Demmel **Statistical Modeling of Feedback Data in an Automatic Tuning System (gzipped ps or pdf)** *Third ACM Workshop on Feedback-Directed and Dynamic Optimization FDDO-3,* Monterey, December 2000

- M. Richardson, J. Bilmes, and C. Diorio **Hidden-Articulator Markov Models: Performance Improvements and Robustness to Noise (gzipped ps or pdf)** *Int. Conf. on Spoken Language Processing,* Beijing, October 2000.

- M. Richardson, J. Bilmes, and C. Diorio **Hidden-Articulatory Markov Models for Speech Recognition (gzipped ps or pdf)** *ISCA ITRW Conference on Automatic Speech Recognition,* Paris, August 2000

- J. Bilmes and K. Kirchhoff **Directed Graphical Models of Classifier Combination: Application to Phone Recognition. (gzipped ps or pdf)** *International Conference on*

Acoustics, Speech, and Signal Processing, Beijing, October 2000.

- D. Ellis and J. Bilmes **Using Mutual Information to Design Feature Combinations (gzipped ps or pdf)** *International Conference on Acoustics, Speech, and Signal Processing*, Beijing, October 2000.
- K. Kirchhoff and J. Bilmes **Combination and joint Training of Acoustic Classifiers for Speech Recognition. (gzipped ps or pdf)** *ISCA ITRW Conference on Automatic Speech Recognition*, Paris, August 2000.
- J. Bilmes. **Dynamic Bayesian Multi-Networks (pdf)** *The 16th Conference on Uncertainty in Artificial Intelligence, Stanford*, July 2000.
- J. Bilmes. **Factored Sparse Inverse Covariance Matrices. (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, June 2000.
- J. Bilmes. **Buried Markov Models for Speech Recognition (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, March 1999.
- K. Kirchhoff and J. Bilmes. **Dynamic Classifier Combination in Hybrid Speech Recognition Systems using Utterance-Level Confidence Values (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, March 1999.
- K. Kirchhoff and J. Bilmes. **Statistical Acoustic Indications of Coarticulation (gzipped ps or pdf)** *14th Int. Cong. of Phonetic Sciences*, August, 1999.
- J. Bilmes. **Data-Driven Extensions to HMM Statistical Dependencies. (gzipped ps or pdf)** *Int. Conf. on Spoken Language Processing,,* Dec 1998.
- J. Bilmes. **Maximum Mutual Information Based Reduction Strategies for Cross-Correlation based Joint Distributional Modeling (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, May 1998.
- J. Bilmes. **Joint Distributional Modeling with Cross-correlation Based Features (gzipped ps or pdf)** *Proc. IEEE Automatic Speech Recognition and Understanding*, Santa Barbara, December 1997.
- J. Bilmes, K. Asanovic, C.-W. Chin, J. Demmel. **Optimizing Matrix Multiply using PHiPAC: a Portable, High-Performance, ANSI C Coding Methodology (gzipped ps or pdf)** *Int. Conf. on Supercomputing,* July 1997.
- J. Bilmes, K. Asanovic, C.-W. Chin, J. Demmel. **Using PHiPAC to Speed Error Back-Propagation Learning (gzipped ps or pdf)** *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, April 1997.
- V. Iyer, J. Bilmes, M. Wright, and D. Wessel **A Novel Representation for Rhythmic Structure (pdf)** *Int. Comp. Mus. Conf,* September 1997.
- J. Bilmes, N. Morgan, S.-L. Wu, and H. Bourlard. **Stochastic Perceptual Speech Models With Durational Dependence (gzipped ps or pdf)** *Intl. Conference on Spoken Language Processing,* Oct 1996.
- J. Bilmes. **Techniques to Foster Drum Machine Expressivity (gzipped ps or pdf)** *International Computer Music Conference*, 1993.
- J. Bilmes. **A Model for Musical Rhythm. (gzipped ps or pdf)** *International Computer Music Conference*, 1992.
- J. Bilmes and D. Anderson. **MOOD: A Concurrent C++-Based Music Language (gzipped ps or pdf)** *International Computer Music Conference*, 1992.
- P. Kohn, J. Bilmes, N. Morgan, and J. Beck. **Software for ANN training on a Ring Array Processor** *NIPS-4*, May 1991.
- H. Schmidt and J. Bilmes. **Exception Handling in pSather.** *Exception Handling*

Workshop, European Conference on Object-Oriented Programming, July 1991.

- J. Bilmes and D. Anderson. **Concurrent Real-Time Music in C++ (gzipped ps or pdf)** *USENIX C++ Conference*, 1989.

Journal Publications and Book Chapters

- C.P.-Chen and J. Bilmes **MVA Processing of Speech Features** , IEEE Transactions on Speech and Audio, (to appear)
- J. Bilmes **Buried Markov Models: A Graphical-Modeling approach to Automatic Speech Recognition** , Computer, Speech and Language, Volume 17, No 2-3, April-July, 2003
- R. Vuduc, J. Demmel, J. Bilmes. **Statistical Models for Empirical Search-Based Performance Tuning.** , International Journal of High Performance Computing Applications, 18(1), pp.65-94, Feb 2004.
- J. Bilmes, **Graphical Models and Automatic Speech Recognition** , in "Mathematical Foundations of Speech and Language Processing", Institute of Mathematical Analysis Volumes in Mathematics Series, Springer-Verlag, 2003.
- M. Richardson, J. Bilmes, and C. Diorio **Hidden-Articulator Markov Models for Speech Recognition** *Speech Communications* 41(2), October 2003.
- J. Bilmes and K. Kirchhoff **Combination and Joint Training of Acoustic Classifiers for Speech Recognition** *Pattern Analysis and Applications* (2003) 6: 201-211.
- N. Morgan, J. Beck, P. Kohn, and J. Bilmes. **Neurocomputing on the RAP** Book chapter in *Parallel Digital Implementations of Neural Networks.* Prentice-Hall, 1993.
- N. Morgan, J. Beck, P. Kohn, J. Bilmes, E. Allman, and J. Beer. **The Ring Array Processor (RAP): A Multiprocessing peripheral for connectionist applications.** *Journal of Parallel and Distributed Computing*, April 1992
- O. Gunther and J. Bilmes. **Tree-Based Access Methods for Spatial Databases: Implementation and Performance Evaluation.** *IEEE Transactions on Knowledge and Data Engineering* , 3(3), 1991.
- O. Gunther and J. Bilmes. **The Implementation of the Cell Tree: Design Alternatives and Performance Evaluation.** *Datenbanksysteme in Buro, Technik und Wissenschaft, Informatik-Fachberichte 204,* , Springer-Verlag, 1989.

Tutorials

- Jeff Bilmes **Graphical Models in Speech and Language Research** , tutorial presented during the 2004 Human Language Technology conference / North American chapter of the Association for Computational Linguistics(HLT/NAACL'04) conference.
- Jeff Bilmes **Graphical Models Research in Audio, Speech, and Language Processing** , tutorial presented during the 2003 Uncertainty in Artificial Intelligence (UAI'03) conference.

Technical Reports and other documents.

- Mukund Narasimhan and Jeff Bilmes **Optimization on Separator Trees** , UWEETR-2004-00018, March 2004 (gzipped postscript or PDF)

- Chia-ping Chen, Jeff Bilmes, Dan P. Ellis **Blind MVA Speech Feature Processing on Aurora 2.0** , UWEETR-2004-00017, March 2004 (gzipped postscript or PDF)
- Chris Bartels and Jeff Bilmes **Elimination is not enough: Non-minimal triangulations for graphical models** , UWEETR-2004-00010, March 2004 (gzipped postscript or PDF)
- Mukund Narasimhan and Jeff Bilmes **Efficient PAC-learning bounded tree-width Graphical Models** , UWEETR-2004-0009, March 2004 (gzipped postscript or PDF)
- Karim Filali, Xiao Li, Jeff Bilmes **Algorithms for Data-Driven ASR Parameter Quantization** (also available in postscript), UWEETR-2002-0010, June, 2003
- Xiao Li, J. Bilmes **Selectively Computing Dynamic Features in the Likelihood Computation of ASR Systems** (also available in postscript), UWEETR-2002-0009, May, 2003
- J. Bilmes, C. Bartels **On Triangulating Dynamic Graphical Models** (also available in postscript), UWEETR-2002-0007, May, 2003
- G. Ji, J. Bilmes, **Necessary Intransitive Likelihood Ratio Classifiers** (also available in postscript), UWEETR-2002-0014, December, 2002
- J. Bilmes, **What HMMs can do.** (also available in postscript), UWEETR-2002-0003, Feb, 2002
- C-.P-. Chen, K. Kirchhof, and J. Bilmes **Towards Simple Methods of Noise Robustness** (also available in postscript), UWEETR-2002-0002, Jan, 2002
- J. Bilmes, et. al. **Discriminatively Structured Graphical Models for Speech Recognition** Johns Hopkins University, CSLP 2001 Summer Workshop Final Report.
- J. Bilmes, **Graphical Models and Automatic Speech Recognition** (also available in postscript), UWEETR-2001-0005, November, 2001
- J. Bilmes, K. Asanovic, C.-W. Chin, J. Demmel. **The PHiPAC v1.0 Matrix-Multiply Distribution** UCB/CSD-98-1020. October, 1998.
- J. Bilmes. **A Gentle Tutorial on the EM Algorithm and its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models** (pdf) other *ICSI-TR-97-021*, 1997.
- J. Bilmes, A. Vahdat, W. Hsu, E.-J. Im. **Empirical Observations of Probabilistic Heuristics for the Clustering Problem** *ICSI-TR097-018*, May 1997.
- J. Bilmes, K. Asanovic, C.-W. Chin, J. Demmel. **PHiPAC: a Portable, High-Performance, C Coding Methodology and its application to matrix multiply.** LAPACK Working Note 111, August 1996.

Technical Documentation

- J. Bilmes. **The Graphical Models Toolkit Documentation** (postscript or pdf).

Ph.D. Thesis

- J. Bilmes. **Natural Statistical Models for Automatic Speech Recognition.** Ph.D. Thesis, Dept. of EECS, CS Division, U.C. Berkeley 1999 (postscript or pdf).

Masters Thesis

- J. Bilmes. **Timing is of the Essence: Perceptual and Computational Techniques for**

Representing, Learning, and Reproducing Expressive Timing in Percussive Rhythm. (gzipped ps or pdf) *Masters Thesis, Massachusetts Institute of Technology, 1993.*